Sensor-Fused Navigation and Manipulation from a Planetary Rover

E. T. Baumgartner[a], P. C. Leger[b], P. S. Schenker[a], and T. L. Huntsberger[c]

[a]Science and Technology Development Section, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109

[b]Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

[c]Intelligent Systems Laboratory, Department of Computer Science, University of South Carolina, Columbia, SC 29208

ABSTRACT

This paper describes the development of advanced rover navigation and manipulation techniques for use by NASA's Sample Return Rover. These techniques include an algorithm for estimating the change in the rover's position and orientation by registering successive range maps from the rover's hazard avoidance stereo camera pair and the fusion of this information with the rover's wheel odometry. This map registration technique is also extended to register range maps to an *a priori* model-based range map for relative rover position and orientation determination. Finally, a technique for the robust and precise positioning of a rover-mounted manipulator using visual feedback from the rover's stereo pair is presented. Experimental results for each of these techniques is documented in this paper.

Keywords: Rover navigation, state estimation, map registration, vision-based manipulation

1. INTRODUCTION

NASA's Jet Propulsion Laboratory (JPL) has recently developed the Sample Return Rover (SRR) shown in Figure 1. This vehicle has been developed as a research and technology demonstration rover for the development of advanced rover navigation concepts.[1] The specific mission requirements that drove the design of this rover included the fast traverse from a lander to a Mars science rover, the acquisition of rock and soil samples collected by the science rover, and the return of these samples to a Mars lander that would launch these samples via a Mars Ascent Vehicle for eventual return to Earth.

Thus, SRR was designed as a four-wheeled, skid-steered rover platform with a split differential and independently-controlled shoulder (or rocker) joints. The computing platform for SRR is a 80486 processor with on-board video framegrabbing, motion control, and analog-to-digital conversion boards. The rover also contains a wireless communications system for linking with a ground control station. SRR carries on-board sensing including a stereo camera pair used for stereo map building and hazard avoidance, a single goal camera for long-range sensing, and an inertial sensing package (three axis gyroscope and accelerometer system). A three degree-of-freedom manipulator arm is also included on the rover that is used primarily for the acquisition of a container which stores the samples (known as a sample cache container). The acquisition of this container is accomplished using a single degree-of-freedom gripping tool located on the manipulator's end-effector.

For the task of rendezvousing and docking with the science rover and its sample cache container, a number of different navigation techniques and algorithms have been developed. These include a RF homing beacon system, visual localization techniques for finding the rover in the goal camera's field-of-view, visual terminal guidance techniques for approaching the science rover, visual feedback techniques for rover state estimation, and, finally, vision-based manipulation algorithms for the robust and reliable acquisition of the sample cache container. This paper describes developments that address the final two technology needs, namely rover state estimation and vision-based manipulation.

Send correspondence to E. T. Baumgartner, E-mail: Eric.T.Baumgartner@jpl.nasa.gov.

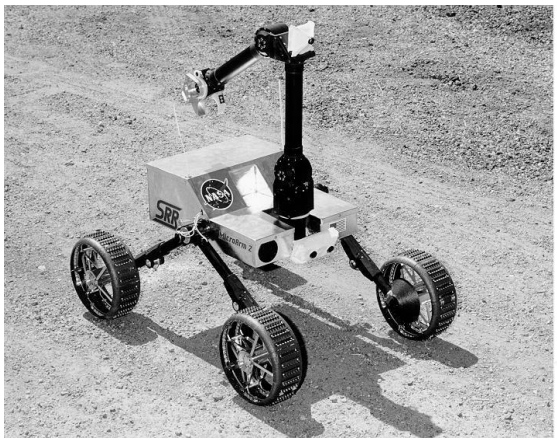

Figure 1. NASA's Sample Return Rover (SRR).

Section 2 describes recent developments and experimental trials for a stereo map registration algorithm that yields a measure of the change in rover's pose (position and orientation) from one stereo frame to another. The adaptation of this map registration technique for accurate terminal positioning of SRR with respect to the science rover's sample cache container is also presented in Section 2. In Section 3, an approach that fuses the results from the map registration technique with rover wheel odometry via an extended Kalman filter framework is discussed. The development of a manipulation technique that utilizes the visual feedback from SRR's stereo camera pair for the robust and reliable positioning of the manipulator's end-effector is documented in Section 4. Finally, some concluding remarks are presented in Section 5.

2. RANGE MAP REGISTRATION

Two key challenges for the SRR vehicle are self-localization and relative pose estimate of the target science rover (or lander) during final approach to the sample cache container. While dead reckoning and beacon following are appropriate for navigating to the immediate vicinity of the target rover, significant errors may occur when executing obstacle avoidance maneuvers. In such cases, 3D registration of successive range maps can be used to estimate rover motion. Terminal guidance of the rover is accomplished by registering an *a priori* model of the goal (rover or lander) to a range map of the goal, thus determining the goal's pose relative to the rover. To address both of these needs, we use a range map registration algorithm based on the Iterative Closest Point (ICP)[2] and Iterative Point Matching (IPM)[3] algorithms. In both cases, range maps are generated from stereo cameras, though the algorithm is applicable to range maps acquired by other means.

2.1. Registration Algorithm

The goal of the registration algorithm is to compute a 3D rotation and translation which aligns one set of data points (S_1) with another (S_2). Our registration algorithm is identical to Zhang's Iterative Point Matching (IPM), with one exception: the surface interpolation of Besl and McKay's Iterative Closest Point (ICP) algorithm is included. We describe these algorithm briefly in this sub-section. Initially, the data points in S_1 are transformed according to estimated rover motion or relative goal pose. Next, the algorithm constructs a k-dimensional (or k-D) tree for the points in S_2, then alternates between two phases: closest-point computation, and estimation of the transformation between corresponding points. The k-D tree enables efficient determination of the closest point in the data set.[4] In the first phase, the algorithm uses the k-D tree to compute the closest point p_2 in the fixed point set (S_2) to each point p_1 in the data set being transformed (S_1). Triangular faces are then created, with each face containing p_2 and two neighbors in the data set. If the distance between p_1 and the closest point on the interpolated triangles is less than a threshold, then p_1 and the closest point are added as a pair to a list of point correspondences. This is where our algorithm differs from IPM: we use ICP's closest-point-on-surface approach, rather than using p_2 as the corresponding point for p_1.

The second phase is performed after the correspondences have been computed. The dual number quaternion method[5] computes a rotation matrix and translation vector which minimize the sum of the squared distances between corresponding points:

$$J = \sum \mathbf{R}\mathbf{x}(i) + \mathbf{T} - \mathbf{y}(i) \tag{1}$$

where $\mathbf{x}(i)$ is the point in S_1 for correspondence i, and $\mathbf{y}(i)$ is its closest point in S_2. Each point in S_1 is transformed by $\mathbf{R}$ and $\mathbf{T}$, and the algorithm continues to iterate the two phases until the magnitudes of $\mathbf{R}$ (represented as a rotation vector) and $\mathbf{T}$ (a translation vector) are each less than a threshold. $\mathbf{R}$ and $\mathbf{T}$ are accumulated at each iteration to compute a final transformation between S_1 and S_2 which can be used for motion estimation or model pose determination.

2.2. Motion Estimation

The registration algorithm, when used with range maps taken from two nearby vehicle poses, can be used to estimate the motion between poses. Since the accuracy of the registration depends on shape variations in the range maps, it is not very useful on flat, featureless terrain. Although this is certainly a limitation, it should be noted that odometry estimates are most reliable when driving over precisely this kind of terrain. On the other hand, in highly-featured terrain where the rover must make frequent turns to avoid obstacles, map registration may be used to refine odometry estimates to increase the accuracy of the rover's knowledge of position and orientation.

When traversing featured terrain, typical rover motions are on the order of 10 cm for travel and 10 degrees for turning. When executing several turns in order to make a large heading change, the rover's skid steering will often cause the wheels to dig in slightly to soft soil, thus making odometry estimates inaccurate. By acquiring stereo maps after each motion and performing map registration, these errors can be reduced to yield an improved motion estimate. We have found that the rotation estimates from map registration are more accurate than those from dead reckoning, but distance estimates are less accurate. This intuitively makes sense, given the uncertainty of each stereo data point: angular uncertainty is relatively low, but range uncertainty is significantly higher. Figure 2 shows the dead-reckoned rover position (dotted line), rover position estimated with map registration (solid line), and measured final (cross). The experiment consisted of 35 motions (15 turns and 20 driving motions), with map registration performed after each movement. Turns were all 10 degrees in magnitude; commanded distances were between 5 and 20 cm. The final error in dead-reckoned pose was 25 cm in position and 10 degrees in heading; map registration yielded a 27 cm position error and 1 degree heading error.

While we have not quantitatively compared our algorithm's performance to that of stereo feature tracking[6] in comparable environments, we feel that there are several promising features of our approach. First, larger motions can be made when using map registration, since the approach does not rely on intensity features appearing similar over successive images. (Note that our system's cameras have approximately 120 degree fields of view and are 20 cm above the ground; a 10 cm translation or 10 degree rotation causes large changes in image feature appearance and location.) However, as mentioned earlier, our approach is not effective on terrain with little shape variation. In this case, stereo feature matching may be more appropriate, since intensity features, rather than shape features, can be used to estimate motion.

Another promising feature of our approach is the ability to register range maps with *a priori* terrain maps, such as might be acquired by high-resolution stereo cameras on a planetary lander. This would be useful only in the vicinity of the lander, but would likely allow more accurate rover localization since registering to a fixed map would tend to eliminate accumulated errors.

2.3. Goal Localization

During the final traverse to the target rover, precise knowledge of the goal pose is crucial. The rover must drive to a location which allows the manipulator arm to retrieve the cache without mechanical interference from the target rover, and precise knowledge of the target rover location can be used to directly compute the location of the cache. This knowledge can be used for cache retrieval, or to verify the cache location computed by other visual methods.

We use several models of the target rover, each acquired using the stereo cameras. The first model is acquired with the target rover position in the desired pose for cache retrieval. Registering this with a range map yields the translation and rotation to the desired retrieval pose. However, since the entire target rover is not within the stereo cameras' fields of view when at the retrieval location, we acquire other stereo models of the target rover from

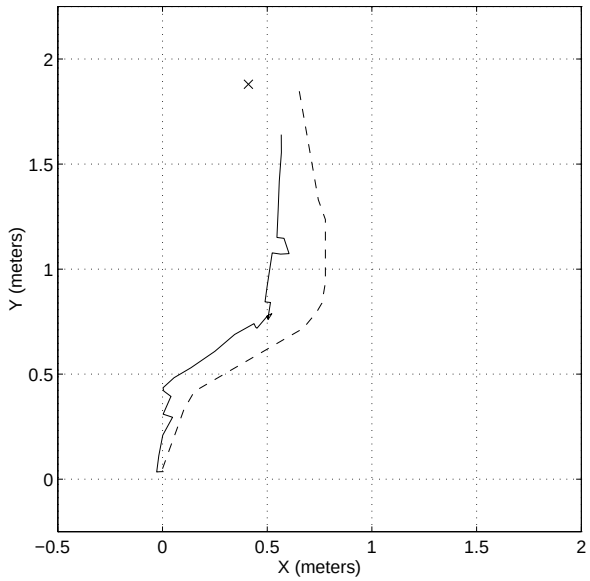

Figure 2. Comparison of dead-reckoned (dotted-line) and map registration (solid line) motion estimates. Final measured rover location is marked by the cross. The final heading was approximately 1 degree; dead reckoning estimated 10 degrees, while map registration estimated 0.3 degrees.

longer ranges so that the entire target rover is visible. These other models may be registered with the first model to transform them to the goal coordinate system.

In our current system, we begin the final approach to the target from approximately 3 meters away. At this point, a goal camera is used to estimate the location of the target. The rover then aligns itself with the target cache and drives to a predetermined standoff distance (currently 1 m) using dead reckoning. The target model is registered with a range map, using the standoff distance as the initial target location. After registration, the registered goal location can be used to guide the rover to the goal. The rover moves a maximum of 10 cm or 10 degrees before acquiring another depth map and refining the goal pose; thus, 10 to 15 motions are required to move to the cache retrieval pose. Typically, the dead-reckoned initial target pose is accurate to 5 cm in range, 15 cm laterally, and 10 degrees in ground-plane rotation. Slightly larger inital range and rotation errors can be tolerated, but larger lateral errors may cause the model to be registered incorrectly. The algorithm is robust to erroneous motion estimates during the final traverse: when using nominal motions of 20 cm and 10 degrees, +/-50% error can be added to each motion estimate without significantly affecting the final localization result. Figure 3 shows a typical localization result at a distance of approximately 50 cm from the target rover. Black points are from the scene data; white points are the model. (Note that this registration was performed with the close-range rover model, which does not include the rear portion of the rover).

When pose errors are within the bounds mentioned above, the rover is able to get within 1 degree and 1-2 cm of the desired position. The computed cache location is accurate to better than 1 cm which is sufficient for cache retrieval. We can likely improve the registration accuracy through more careful calibration and using higher-resolution stereo maps during the final motions. We anticipate being able to use the same model registration algorithm for returning to and docking with a lander to hand off the sample cache for return to earth.

3. SENSOR FUSION

As described in the previous section, the map registration technique results in accurate estimates of the rover's change in orientation while the estimates of the rover's change in translation are not as accurate. In direct contrast to the map registration results, using wheel odometry alone (i.e. dead-reckoning), estimates of the rover's change in translation are quite accurate while estimates of a change in rover orientation are not very accurate. This inaccuracy

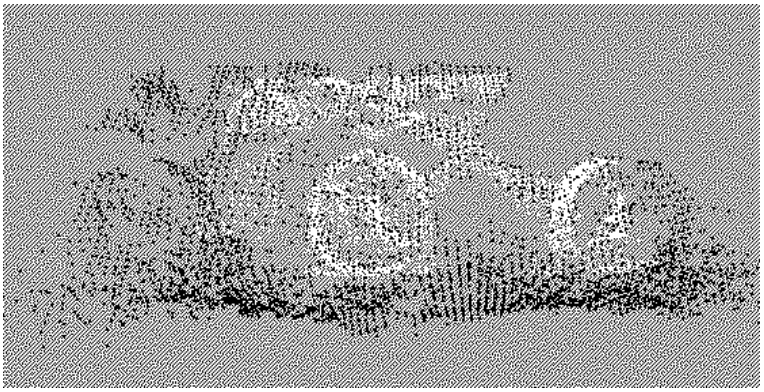

Figure 3. Typical model registration result. The model (white) has been registered to the scene data (black), which was acquired at distance of approximately 50 cm from the target rover.

is primarily due to the use of skid-steering as the means for turning the rover and the impact of wheel slippage on the ability to accurately skid-steer the rover to a particular orientation.

Therefore, the development of an approach that fuses together these two independent measures of the rover's pose by taking advantage of knowledge of the confidence placed on each measurement source is required. Such a sensor fusion framework has been developed through the use an extended Kalman filter technique first proposed by Baumgartner and Skaar[7] for rover state estimation using vision and wheel odometry. This technique has been documented in detail in a paper by Hoffman *et al*[8] and is discussed briefly here.

The map registration technique presented in the previous section produces a measure of the change in translation and orientation of the rover. However, this measurement will generally be in error to some degree due to camera calibration errors, finite stereo correlation precision, etc. Likewise, wheel odometry will be in error throughout a rover traverse due to wheel slippage, incorrectly-modeled wheel kinematics, etc. Therefore, the sensor fusion approach utilizes an extended Kalman filter to propagate the rover kinematic state equations between rover map registrations using the following state equation:

$$\begin{aligned} \frac{d\mathbf{x}(\alpha)}{d\alpha} &= \begin{bmatrix} \frac{dX(\alpha)}{d\alpha} \\ \frac{dY(\alpha)}{d\alpha} \\ \frac{d\phi(\alpha)}{d\alpha} \end{bmatrix} = \begin{bmatrix} R\cos\phi(\alpha) \\ R\sin\phi(\alpha) \\ \frac{R}{B}u \end{bmatrix} + \mathbf{w}(\alpha) \\ &= \mathbf{f}(\mathbf{x}(\alpha), u) + \mathbf{w}(\alpha) \end{aligned} \tag{2}$$

where $\mathbf{x} = [X\ Y\ \phi]^T$ is the state vector of vehicle (planar position and orientation), the independent variable α represents the average rover wheel rotation, u is the control variable, and $\mathbf{w}(\alpha)$ is known as the process noise. The process noise is assumed to be a Gaussian-distributed, white-noise process with covariance matrix, $\mathbf{Q}$. The map registration results are then represented by the following measurement equation:

$$\begin{aligned} \mathbf{z}(\alpha_i) &= \begin{bmatrix} X(\alpha_{i-1}) + \Delta X(\alpha_i) \\ Y(\alpha_{i-1}) + \Delta Y(\alpha_i) \\ \phi(\alpha_{i-1}) + \Delta\phi(\alpha_i) \end{bmatrix} + \mathbf{v}(\alpha_i) \\ &= \mathbf{h}(\mathbf{x}(\alpha_i)) + \mathbf{v}(\alpha_i) \end{aligned} \tag{3}$$

where ΔX, ΔY, and $\Delta\phi$ are computed directly from $\mathbf{R}$ and $\mathbf{T}$ as detailed in Section 2. Also, in equation 3, α_i represents the value of α at which time a stereo pair is acquired and the map registration algorithm is completed and α_{i-1} is the value of α at which the previous vision-based update to the Kalman filter was computed. The measurement process is assumed to be a Gaussian-distributed, white-noise process with covariance matrix, $\mathbf{R}$. These propagated state estimates are then fused together with the measurement via the extended Kalman filter update equations as presented in Baumgartner and Skaar[7] and in Hoffman *et al*.[8] This update attempts to reduce the

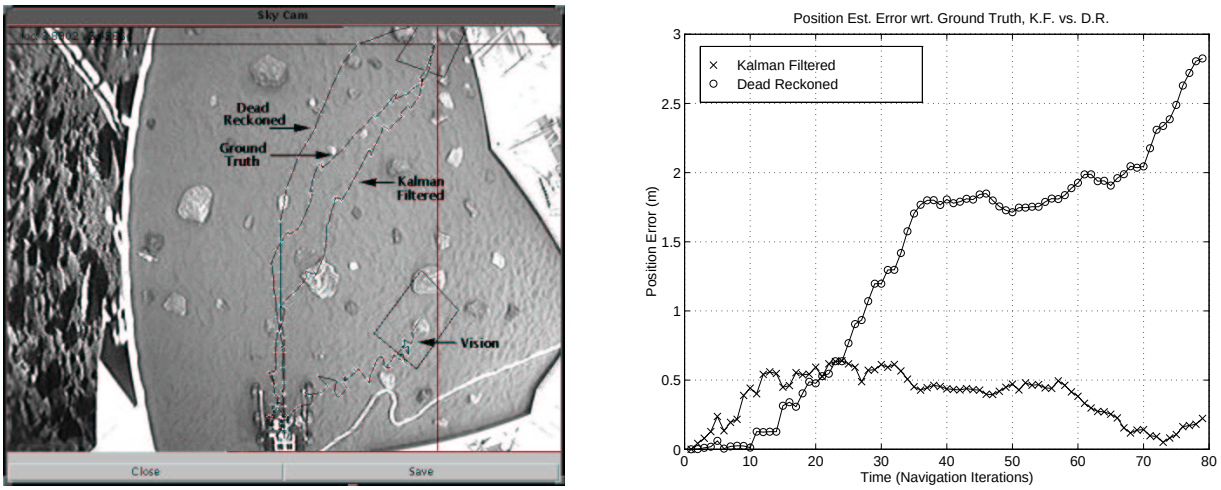

Figure 4. Estimation results for the LSR vehicle showing dead-reckoned, map registration, and extended Kalman filter state estimates.

estimation error covariance matrix, $\mathbf{P}(\alpha)$, through knowledge of the covariances associated with the process and measurement noise.

This state estimation approach has been experimentally demonstrated using NASA's Lightweight Survivable Rover (LSR)[9] platform which is a six-wheeled, skid-steered rover with a stereo hazard avoidance camera system similar to the one used on the SRR platform. Extensive testing of the EKF state estimator within JPL's Planetary Robotics Laboratory which contains a rover pit filled with a soft soil material and is populated with a wide variety of large and small rocks. This terrain type represents a difficult rover traverse environment due to the soft soil and resulting wheel slippage and the numerous rock obstacles. Figure 4 shows a results from the state estimator during an obstacle avoidance traverse across the rover pit. Over this six meter traverse, the EKF state estimator produced a final state estimate error of only 22 centimeters in translation with a maximum estimate error of 64 centimeters during the traverse. Using wheel odometry alone to determine the rover state knowledge, the final state estimation error during this same run was nearly 3 meters in translation. Therefore, this work has shown that the sensor fusion of wheel odometry with the vision-based map registration technique via the EKF framework can have a profound impact on the quality of the rover's state estimates.

4. HYBRID IMAGE-PLANE/STEREO (HIPS) MANIPULATION

One of the most important tasks of the Sample Return Rover is the robust and reliable acquisition of the sample cache container. This is accomplished using the rover-mounted 3 degree-of-freedom MicroArm-2 manipulator and its single degree-of-freedom gripping end-effector. The baseline approach for this manipulation task is the localization of the sample container with respect to the manipulator arm via the hazard-avoidance stereo cameras located on SRR. Once the sample container has been identified in the image plane of one of the stereo cameras, the 3D range to the container and the container's handle can be determined via stereo correlation and triangulation. From this 3D range information, the joint rotations that take the manipulator to the appropriate location in 3D space are determined using the arm's inverse kinematics.

The difficulty associated with this approach is that sources of error tend to accumulate in the stereo system and the manipulator kinematics that ultimately reduce the ability to accurately place the manipulator's end-effector at the desired location. Such error sources include stereo calibration and stereo ranging errors that can yield a range uncertainty of 1.5 cm at 0.75 meters from the stereo pair used on SRR. Errors are also associated with accurate knowledge of the transformation between the stereo reference frame to the base reference frame of the manipulator as well as kinematic errors such as link length uncertainties and joint position uncertainties. Finally, degradations and/or changes in the system configuration due to environmental factors such as launch and landing vibrations

and thermal expansion will also effect the ability to accurately position a rover-mounted manipulator at a target of interest.

Therefore, we have begun the development of a vision-based manipulation technique known as Hybrid Image-Plane/Stereo (HIPS) manipulation that addresses the issues raised above with respect to pointing precision. This technique is an extension of the work by Skaar *et al*[10,11] which has been come to be known as Camera Space Manipulation (CSM). In the CSM approach, widely separated cameras are utilized to determine the direct relationship between the joint rotations of a manipulator and the image-plane (or camera-space) appearance of a "cue" located on the manipulator's end-effector. The inverse problem is then solved to determine the joint rotations that drive the end-effector to a particular image-plane target location in the participating cameras without regard to any physical reference frame. This method achieves excellent terminal pointing precision (sub-millimeter) when the participating cameras are spaced nearly perpendicular to each other. Unfortunately, widely-spaced cameras on a rover platform is difficult to achieve due to the finite size of the rover and due to the use of existing rover cameras which are configured as stereo pairs for rover navigation needs (i.e. the addition of manipulator-specific cameras is unattractive due to mass, volume, power, and complexity issues).

The HIPS manipulation technique makes use of the basic idea underlying the CSM method – the generation of camera models using visual sensing of the manipulator's end-effector and the use of these camera models to drive the end-effector to a target location without regard to any "real" physical reference frame. In HIPS, the calibration of each of the stereo cameras is conducted using a visual "cue" located on the manipulator's end-effector. This is implemented by positioning the end-effector throughout a set of joint rotations that span workspace of the manipulator as well as the image plane of each of the stereo cameras. The 3D location of the end-effector at each joint rotation set is computed using the nominal arm kinematics. Once this data set has been collected, the parameters in the stereo camera models are determined. It is important to note that the camera model developed may or may not have any relationship to the camera model generated using an accurate calibration fixture, however, the camera models generated do accurately reflect the relationship between the position of the end-effector (known via the manipulator's forward kinematics) and image-plane appearance of the end-effector in each of the stereo cameras.

Once these manipulator-generated camera models have been determined as described above, a target in the image plane is selected and these camera models are utilized to solve for the 3D location of this target via stereo correlation and triangulation. Again, the actual 3D location of this target may be quite different from the 3D location determined by the manipulator-generated camera models. The important point is that the 3D location determined by the manipulator-generated camera models is accurate with respect to the manipulator which is the means by which the end-effector positioning will be realized. Thus, with the 3D location of the target known, the nominal inverse kinematics are used to solve for the joint rotations of the manipulator that places the end-effector at the desired target location.

This approach addresses the systematic errors that are present in the baseline technique (separate camera calibration and manipulator kinematics), however, stochastic errors that occur due to finite image-plane cue detection precision, inaccurate knowledge of joint angles, etc. are not necessarily accounted for using the HIPS manipulation technique. Therefore, the manipulator-generated camera models are updated throughout the trajectory that takes the end-effector to the target by identifying the image-plane location the visual cue on the end-effector and by using this additional information to re-compute the parameters in the camera model. In this way, the manipulator-generated camera models are refined such that they are quite accurate in the vicinity of the terminal target location.

To test the HIPS manipulation technique, two studies were performed. The first is a simulation study that introduces a significant error into the arm kinematics (a 1.9 cm change in one of the arm's nominal link length) without giving knowledge to or modifying the nominal kinematic model of the manipulator. A "truth" camera model is generated based on the modified kinematic model and used to determine the image plane appearance of a cue on the modified manipulator in each of the stereo cameras. The stereo camera calibrations are then generated using the image-plane locations of the cue using the "truth" model. The 3D location of the cue is computed using the nominal (and in error) arm kinematics. The results of this camera calibration are shown in Figure 5 for each of the stereo cameras. These figures indicate that the systematic error in the link length change is accounted for in the manipulator-generated camera models. Finally, an image-plane target location is chosen and the 3D range to this target is determined using the manipulator-generated camera models. The inverse kinematics using the nominal (and, again, in error) link length are utilized to solve for the joint rotations that take the end-effector to the target location. Since this is a simulation, the actual 3D location of the target is known and when compared to the 3D

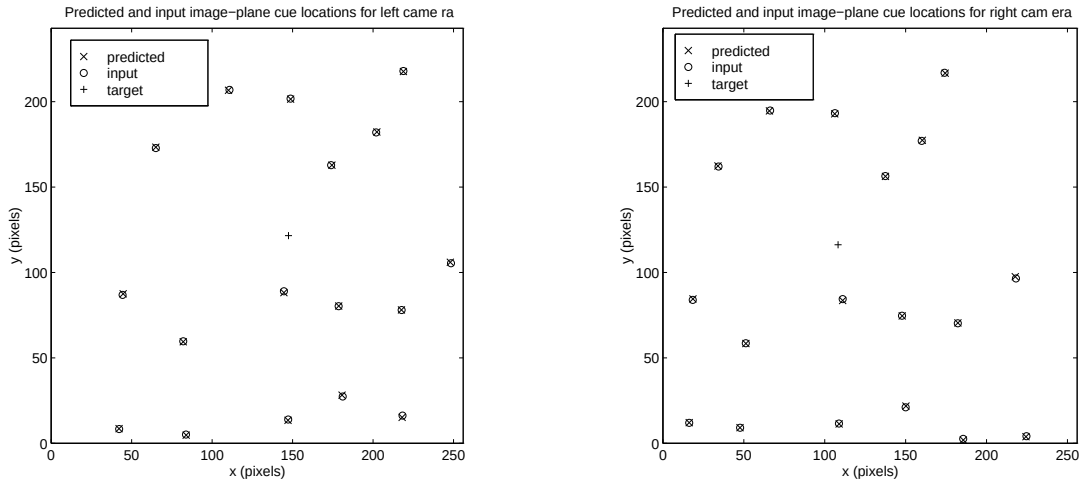

Figure 5. HIPS simulation results with significant kinematic modeling error, left and right image-planes.

location of the end-effector that is realized using the joint rotations calculated with HIPS, a final positioning error of 2.5 mm is realized. Had the nominal arm kinematics been used along with the nominal stereo camera calibration (as in the baseline approach), the final positioning error would be approximately 1.9 cm. Thus, the HIPS manipulation technique reduced the systematic errors in this study by almost a factor of 10.

Similarly, a study was conducted using the manipulator, MicroArm-2, on the SRR vehicle. For this study, the manipulator's end-effector is moved throughout its workspace within the field-of-view of the stereo hazard avoidance cameras on SRR. The image-plane location of a point on the end-effector is identified and the 3D location of this point is determined using the manipulator's kinematic model. Using this information, the manipulator-generated camera models for each camera in the stereo pair is computed. Then, an image-plane target location is specified and the 3D range to this target is determined using the manipulator-generated camera models. The inverse kinematics of the manipulator are then used to determine the joint rotations that take the end-effector to this target location. Since the 3D location of this target is known via the arm's forward kinematics, the 3D location of the end-effector as calculated using the joint rotations returned from the HIPS technique can be utilized to determine the final positioning error. In this case, the final positioning error was 7.8 mm. It is important to note that the manipulator-generated camera models were not updated throughout the end-effector trajectory to the target location. However, the resulting positioning error of 7.8 mm is a significantly improvement over the positioning error that is realized using the baseline approach (typically on the order of 1.5 - 2 cm depending on the range to the target). The results of the manipulator-generated camera calibration and the target location are shown in Figure 6 for each of SRR's stereo camera image-planes.

Future work for the development of the HIPS manipulation technique includes the full-scale implementation on SRR with automatic cue detection. This will allow the *in situ* calibration of the manipulator-generated camera models so that precision end-effector positioning can be realized throughout the rover's mission life despite changes or degradations to the hardware. We will also detail and characterize the performance of the method through numerous experimental trials. Currently, the method is targeted for a 3 degree-of-freedom manipulator, so only the 3D positioning problem has been addressed. The method will be further generalized to handle the 6D problem (3D positioning and orientation) for use on rover-mounted manipulators with higher degrees-of-freedom.

5. CONCLUSIONS

This paper has described advanced rover navigation and manipulation technology developments. Experimental results generated using JPL's Sample Return Rover and Lightweight Survivable Rover have demonstrated the ability to accurately estimate the rover's position and orientation via the sensor fusion of wheel odometry and map registration results. We also show that 3D model registration for target localization allows for the precise terminal guidance of the SRR for sample cache retrieval. This paper has also introduced the Hybrid Image-Plane/Stereo Manipulation

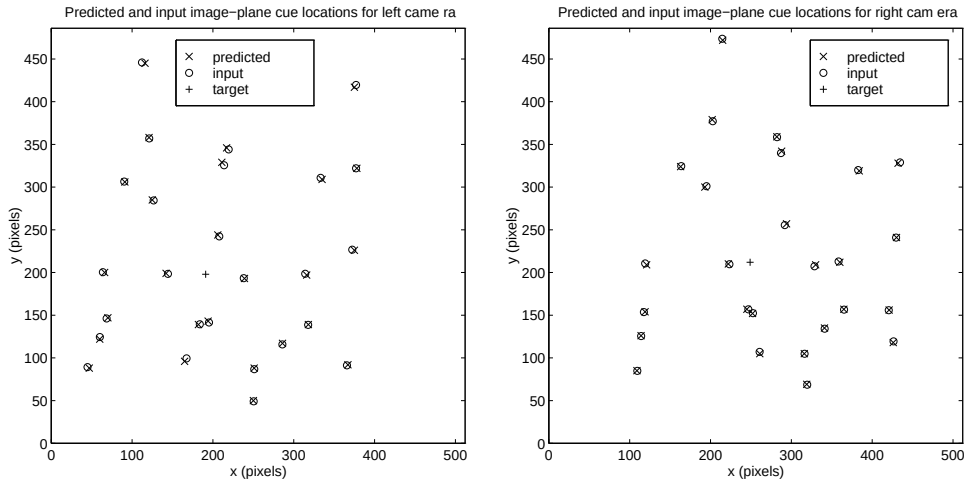

Figure 6. HIPS results for SRR's MicroArm-2 manipulator, left and right image-planes.

technique for the robust and accurate placement of a rover-mounted robot arm using visual feedback from the rover's stereo camera pair. Both of these techniques will be refined and tested extensively through many experimental trials so that the resulting algorithms will represent a significant advance in rover technology.

ACKNOWLEDGMENTS

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

1. P. S. Schenker, E. T. Baumgartner, R. A. Lindemann, H. Aghazarian, A. J. Ganino, G. S. Hickey, D. Q. Zhu, L. H. Matthies, B. H. Hoffman, and T. L. Huntsberger, "New planetary rovers for long-range mars science and sample return," *to appear in Proc. SPIE* **3522**, 1998.
2. P. Besl and N. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **14**(2), pp. 239–256, 1992.
3. Z. Zhang, "Iterative point matching for registration of free-form curves and surfaces," *International Journal of Computer Vision* **13**(2), pp. 119–152, 1994.
4. F. Preparata and M. Shamos, *Computational Geometry, An Introduction*, Springer, Berling, Heidelberg, New York, 1986.
5. M. W. Walker, L. Shao, and R. A. Volz, "Estimating 3-d location parameters using dual number quaternions," *CVGIP: Image Understanding* **54**(3), pp. 358–367.
6. L. Matthies, *Dynamic stereo vision*, Ph. D. Thesis, Canegie Mellon University, 1987.
7. E. T. Baumgartner and S. B. Skarr, "An autonomous vision-based mobile robot," *IEEE Transactions on Automatic Control* **39**(3), March 1994.
8. B. D. Hoffman, E. T. Baumgartner, T. Hutsberger, and P. S. Schenker, "Improved rover state estimation in challenging terrain," submitted for review to *Autonomous Robots* , February 1992.
9. P. S. Schenker, L. F. Sword, A. J. Ganino, D. B. Bickler, G. S. Hickey, D. K. Brown, E. T. Baumgartner, L. H. Matthies, B. H. Wilcox, T. Balch, H. Aghazarian, and M. S. Garrett, "Lightweight rovers for mars science exploration and sample return," *Proc. SPIE* **3208**, pp. 24–36, 1997.
10. S. B. Skaar, W. H. Brockman, and W. S. Jang, "Three-dimensional camera space manipulation," *The International Journal of Robotics Research* **9**(4), pp. 22–39, August 1990.
11. W. Z. Chen, U. Korde, and S. B. Skaar, "Position-control experiments using vision," *International Journal of Robotics Research* **13**(3), pp. 199–208, June 1994.